

PATENT

10830.099.NPUS00

APPLICATION FOR UNITED STATES LETTERS PATENT

for

**MAINTENANCE OF A FILE VERSION SET INCLUDING READ-ONLY
AND READ-WRITE SNAPSHOT COPIES OF A PRODUCTION FILE**

By

Peter Bixby

Sachin Mullick

Jiannan Zheng

Xiaoye Jiang

Sorin Faibish

Express Mail Mailing Label No. _____

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to a file server maintaining snapshot copies of a read-write file by sharing file blocks and performing a copy-on-write to a newly allocated block when writing to a shared block.

2. Description of the Related Art

A technique known as file versioning maintains read-only versions of a read-write production file by sharing file blocks between the production file and the read-only versions, and performing a copy-on-write to a newly allocated block for the production file when writing to a shared block. Such a file system is described in Chutani, Sailesh, et al., "The Episode File System," Carnegie Mellon University IT Center, Pittsburgh, PA, June 1991, incorporated herein by reference. Each read-only version is a snapshot of the production file at a respective point in time. Read-only versions can be used for on-line data backup and data mining tasks.

In a copy-on-write file versioning method, the read-only version initially includes only a copy of the inode of the production file. Therefore the read-only version initially shares all of the data blocks as well as any indirect blocks of the production file. When the production file is modified, new blocks are allocated and linked to the production file inode to save the new data, and the original data blocks are retained and linked to the inode of the read-only version. The result is that disk space is saved by only saving the difference between two consecutive versions. If the production file becomes corrupted during a system crash, then typically the most recent read-only version is copied over to

1 the production file in a recovery operation. In this case, there is a loss of the data that
2 was written to the production file since the creation of the most recent read-only version.

3 Users are becoming less tolerant of delays in accessing their data, and even less
4 tolerant of corruption of their data. Therefore, there has been a continuing interest in
5 improving data availability and the effectiveness of recovery procedures. For example,
6 after recovery, the integrity of the recovered file is checked, and if a defect is found, an
7 attempt is made to correct it. In addition, it is often possible to recover some of the data
8 that was written to the production file since the creation of the latest read-only version,
9 for example, by replay of a log from an application program.

10

11

SUMMARY OF THE INVENTION

12

13

14

15

16

17

18

19

20

In accordance with one aspect, the invention provides a file server. The file
server includes storage containing a file system, and a processor coupled to the storage
for accessing the file system. The file system includes a production file, read-only
snapshot copies of the production file, and at least one read-write snapshot copy of the
production file. The production file and the snapshot copies of the production file are
organized as a version set including a node for the production file and a node for each
snapshot copy of the production file, and a set of file blocks including direct and indirect
blocks that are shared among the production file and the snapshot copies of the
production file.

21

22

23

In accordance with another aspect, the invention provides a file server. The file
server includes storage containing a file system, and a processor coupled to the storage
for accessing the file system. The file system includes a production file, read-only

1 snapshot copies of the production file, and at least one read-write snapshot copy of the
2 production file. The production file and the snapshot copies of the production file are
3 organized as a version set including a node for the production file and a node for each
4 snapshot copy of the production file, and a set of file blocks including direct and indirect
5 blocks that are shared among the production file and the snapshot copies of the
6 production file. The file server further includes means for creating new read-only
7 snapshot copies of the production file, means for creating new read-write snapshot copies
8 of the production file, means for deleting a specified snapshot copy of the production file
9 from the version set, means for restoring the production file with a specified snapshot
10 copy of the production file, means for refreshing a specified snapshot copy of the
11 production file, and means for naming the files in the version set.

12 In accordance with another aspect, the invention provides a file server. The file
13 server includes storage containing a file system, and a processor coupled to the storage
14 for accessing the file system. The file system includes a production file, and read-only
15 snapshot copies of the production file. The production file and the read-only snapshot
16 copies of the production file are organized as a version set including a node for the
17 production file, a node for each read-only snapshot copy of the production file, and a set
18 of file blocks including direct and indirect blocks that are shared among the production
19 file and the read-only snapshot copies of the production file. The file server is
20 programmed to maintain for each block in each version of the production file an
21 indication of whether or not the version of the production file is an oldest version of the
22 production file including an identical version of the block. The file server is
23 programmed to delete a read-only snapshot copy of the production file, and when

1 deleting the read-only snapshot copy of the production file, to keep each block for which
2 the read-only snapshot copy is not indicated as being an oldest version of the production
3 file including an identical version of the block.

4 In accordance with another aspect, the invention provides a file server. The file
5 server includes storage containing a file system, and a processor coupled to the storage
6 for accessing the file system. The file system includes a production file, and snapshot
7 copies of the production file. The production file and the snapshot copies of the
8 production file are organized as a version set including a node for the production file, a
9 node for each snapshot copy of the production file, and a set of file blocks including
10 direct and indirect blocks that are shared among the production file and the snapshot
11 copies of the production file. The file server is programmed for responding to a request
12 to create a read-only snapshot copy of the production file by reserving for the production
13 file a number of free file blocks of at least the number of blocks in the production file.

14 In accordance with another aspect, the invention provides a file server. The file
15 server includes storage containing a file system, and a processor coupled to the storage
16 for accessing the file system. The file system includes a production file, and snapshot
17 copies of the production file. The production file and the snapshot copies of the
18 production file are organized as a version set including a node for the production file, a
19 node for each snapshot copy of the production file, and a set of file blocks including
20 direct and indirect blocks that are shared among the production file and the snapshot
21 copies of the production file. The file server is programmed for restoring the production
22 file with a specified snapshot copy of the production file by responding to a request to
23 prepare to restore the production file by preparing to restore the production file and

1 reporting whether or not preparation is successful, and then responding a request to
2 commit the preparation by restoring the production file with the specified snapshot copy
3 of the production file.

4 In accordance with another aspect, the invention provides a file server. The file
5 server includes storage containing a file system, and a processor coupled to the storage
6 for accessing the file system. The file system includes a production file, and snapshot
7 copies of the production file. The production file and the snapshot copies of the
8 production file are organized as a version set including a node for the production file, a
9 node for each snapshot copy of the production file, and a set of file blocks including
10 direct and indirect blocks that are shared among the production file and the snapshot
11 copies of the production file. The file server is programmed for refreshing a specified
12 snapshot copy of the production file by creating a new node in the version set, copying
13 contents of the node of the specified snapshot copy into the new node so that the new
14 node references blocks of the specified snapshot copy, using the node of the specified
15 snapshot copy to create a new snapshot copy of the production file by copying contents
16 of the node of the production file into the node of the specified snapshot copy, and
17 performing a file deletion upon the new node.

18 In accordance with another aspect, the invention provides a method of operating a
19 file server. The file server includes storage containing a file system, and a processor
20 coupled to the storage for accessing the file system. The file system includes a
21 production file, and read-only snapshot copies of the production file. The production file
22 and the read-only snapshot copies of the production file are organized as a version set
23 including a node for the production file, a node for each read-only snapshot copy of the

1 production file, and a set of file blocks including direct and indirect blocks that are shared
2 among the production file and the read-only snapshot copies of the production file. The
3 method includes maintaining for each block in each version of the production file an
4 indication of whether or not the version of the production file is an oldest version of the
5 production file including an identical version of the block. The method further includes
6 deleting a read-only snapshot copy of the production file, wherein the deleting of the
7 read-only snapshot copy of the production file includes keeping each block for which the
8 read-only snapshot copy is not indicated as being an oldest version of the production file
9 including an identical version of the block.

10 In accordance with another aspect, there is provided a method of operating a file
11 server. The file server includes storage containing a file system, and a processor coupled
12 to the storage for accessing the file system. The file system includes a production file,
13 and snapshot copies of the production file. The production file and the snapshot copies of
14 the production file are organized as a version set including a node for the production file,
15 a node for each snapshot copy of the production file, and a set of file blocks including
16 direct and indirect blocks that are shared among the production file and the snapshot
17 copies of the production file. The method includes the file server responding to a request
18 to create a read-only snapshot copy of the production file, and when responding to the
19 request to create a read-only snapshot copy of the production file, reserving for the
20 production file a number of free file blocks of at least the number of blocks in the
21 production file.

22 In accordance with yet another aspect, the invention provides a method of
23 operating a file server. The file server includes storage containing a file system, and a

1 processor coupled to the storage for accessing the file system. The file system includes a
2 production file, and snapshot copies of the production file. The production file and the
3 snapshot copies of the production file are organized as a version set including a node for
4 the production file, a node for each snapshot copy of the production file, and a set of file
5 blocks including direct and indirect blocks that are shared among the production file and
6 the snapshot copies of the production file. The method includes the file server restoring
7 the production file with a specified snapshot copy of the production file by responding to
8 a request to prepare to restore the production file by preparing to restore the production
9 file and reporting whether or not preparation is successful, and then responding a request
10 to commit the preparation by restoring the production file with the specified snapshot
11 copy of the production file.

12 In accordance with a final aspect, there is provided a method of operating a file
13 server. The file server includes storage containing a file system, and a processor coupled
14 to the storage for accessing the file system. The file system includes a production file,
15 and snapshot copies of the production file. The production file and the snapshot copies
16 of the production file are organized as a version set including a node for the production
17 file, a node for each snapshot copy of the production file, and a set of file blocks
18 including direct and indirect blocks that are shared among the production file and the
19 snapshot copies of the production file. The method includes the file server refreshing a
20 specified snapshot copy of the production file by creating a new node in the version set,
21 copying contents of the node of the specified snapshot copy into the new node so that the
22 new node references blocks of the specified snapshot copy, using the node of the
23 specified snapshot copy to create a new snapshot copy of the production file by copying

1 contents of the node of the production file into the node of the specified snapshot copy,
2 and performing a file deletion upon the new node.

3

4 **BRIEF DESCRIPTION OF THE DRAWINGS**

5 Other objects and advantages of the invention will become apparent upon reading
6 the following detailed description with reference to the accompanying drawings wherein:

7 FIG. 1 is a block diagram of a data processing system including multiple clients
8 and a network file server;

9 FIG. 2 is a block diagram showing further details of the network file server in the
10 data processing system of FIG. 1;

11 FIG. 3 is a block diagram of various read and write interfaces in a Unix-based file
12 system layer (UxFS) in the network file server of FIG. 2;

13 FIG. 4 shows various file system data structures associated with a file in the
14 network file server of FIG. 2;

15 FIGS. 5 and 6 comprise a flowchart of programming in the Common File System
16 (CFS) layer in the network file server for handling a write request from a client;

17 FIG. 7 is a timing diagram showing multiple read and write operations pipelined
18 into parallel streams in the Common File System (CFS) layer in the network file server
19 for handling concurrent write requests from a client;

20 FIG. 8 shows multiple processors for processing the pipelined read and write
21 operations in the network file server;

22 FIG. 9 is a flowchart of programming in the Common File System (CFS) layer in
23 the network file server for handling a read request from a client;

1 FIG. 10 is a flowchart of programming in the Common File System (CFS) layer
2 in the network file server for handling concurrent read and write requests from a client;

3 FIG. 11 is a flowchart of a write thread in the UxFS layer of the network file
4 server;

5 FIG. 12 is a more detailed flowchart of steps in the write thread for committing
6 preallocated metadata;

7 FIG. 13 is a block diagram of a partial block write during a copy-on-write
8 operation;

9 FIG. 14 is a block diagram of a read-write file as maintained by the UxFS layer;

10 FIG. 15 is a block diagram of the read-write file of FIG. 14 after creation of a
11 read-only snapshot copy of the read-write file;

12 FIG. 16 is a block diagram of the read-write file of FIG. 15 after a copy-on-write
13 operation upon a data block and two indirect blocks between the data block and the inode
14 of the read-write file;

15 FIG. 17 is a flowchart of steps in a write thread for performing the partial block
16 write operation of FIG. 13;

17 FIG. 18 shows a flowchart of steps in a write thread for allocating file blocks
18 when writing to a file having read-only snapshots;

19 FIG. 19 is a block diagram of a file version set including read-only and read-write
20 snapshot copies of a production file;

21 FIG. 20 is a flowchart of a procedure for creating a new production file;

22 FIG. 21 is a block diagram of a conventional inode of a file;

23 FIG. 22 is a block diagram of an inode in the file version set of FIG. 19;

1 FIG. 23 is a block diagram showing linkages between the inodes in the file
2 version set of FIG. 19;

3 FIG. 24 is a flowchart of a procedure for creating a read-only snapshot copy in the
4 file version set of FIG. 19;

5 FIG. 25 is a flowchart of a procedure for creating a read-write branch in the file
6 version set of FIG. 19;

7 FIG. 26 is a flowchart of a procedure for deleting a read-only version in the file
8 version set of FIG. 19;

9 FIGS. 27-28 comprise a flowchart of a procedure for reserving file blocks for
10 read-write files in order to ensure that the sharing of file blocks among the files in the
11 version set of FIG. 19 is not likely to result in a shortage of file blocks when writing to
12 the read-write files;

13 FIG. 29 is a state diagram for the process of restoring a production file with a
14 read-only version;

15 FIG. 30 is a flowchart of a procedure for preparing for the restoration of the
16 production file with a read-only version;

17 FIG. 31 is a flowchart of a procedure for aborting the restoration of the production
18 file with a read-only version;

19 FIG. 32 is a flowchart of a procedure for committing the restoration of the
20 production file with a read-only version;

21 FIG. 33 is a flowchart of a procedure for refreshing a read-only version; and

22 FIGS. 34 and 35 comprise a flowchart of a procedure for parsing a file name for a
23 file in the version set of FIG. 19.

1 While the invention is susceptible to various modifications and alternative forms,
2 specific embodiments thereof have been shown in the drawings and will be described in
3 detail. It should be understood, however, that it is not intended to limit the invention to
4 the particular forms shown, but on the contrary, the intention is to cover all
5 modifications, equivalents, and alternatives falling within the scope of the invention as
6 defined by the appended claims.

7

8 **DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS**

9 I. The preferred network file server.

10 FIG. 1 shows an Internet Protocol (IP) network 20 including a network file server
11 21 and multiple clients 23, 24, 25. The network file server 21, for example, has multiple
12 data mover computers 26, 27, 28 for moving data between the IP network 20 and a
13 cached disk array 29. The network file server 21 also has a control station 30 connected
14 via a dedicated dual-redundant data link 31 among the data movers for configuring the
15 data movers and the cached disk array 29.

16 Further details regarding the network file server 21 are found in Vahalia et al.,
17 U.S. Patent 5,893,140, incorporated herein by reference, and Xu et al., U.S. Patent
18 6,324,581, issued Nov. 27, 2001, incorporated herein by reference. The network file
19 server 21 is managed as a dedicated network appliance, integrated with popular network
20 operating systems in a way, which, other than its superior performance, is transparent to
21 the end user. The clustering of the data movers 26, 27, 28 as a front end to the cached
22 disk array 29 provides parallelism and scalability. Each of the data movers 26, 27, 28 is a
23 high-end commodity computer, providing the highest performance appropriate for a data

1 mover at the lowest cost. The data mover computers 26, 27, 28 may communicate with
2 the other network devices using standard file access protocols such as the Network File
3 System (NFS) or the Common Internet File System (CIFS) protocols, but the data mover
4 computers do not necessarily employ standard operating systems. For example, the
5 network file server 21 is programmed with a Unix-based file system that has been
6 adapted for rapid file access and streaming of data between the cached disk array 29 and
7 the data network 20 by any one of the data mover computers 26, 27, 28.

8 FIG. 2 shows software modules in the data mover 26 introduced in FIG. 1. The
9 data mover 26 has a Network File System (NFS) module 41 for supporting
10 communication among the clients and data movers of FIG. 1 over the IP network 20
11 using the NFS file access protocol, and a Common Internet File System (CIFS) module
12 42 for supporting communication over the IP network using the CIFS file access
13 protocol. The NFS module 41 and the CIFS module 42 are layered over a Common File
14 System (CFS) module 43, and the CFS module is layered over a Universal File System
15 (UxFS) module 44. The UxFS module supports a UNIX-based file system, and the CFS
16 module 43 provides higher-level functions common to NFS and CIFS.

17 The UxFS module accesses data organized into logical volumes defined by a
18 module 45. Each logical volume maps to contiguous logical storage addresses in the
19 cached disk array 29. The module 45 is layered over a SCSI driver 46 and a Fibre-
20 channel protocol (FCP) driver 47. The data mover 26 sends storage access requests
21 through a host bus adapter 48 using the SCSI protocol, the iSCSI protocol, or the Fibre-
22 Channel protocol, depending on the physical link between the data mover 26 and the
23 cached disk array 29.

1 A network interface card 49 in the data mover 26 receives IP data packets from
2 the IP network 20. A TCP/IP module 50 decodes data from the IP data packets for the
3 TCP connection and stores the data in message buffers 53. For example, the UxFS layer
4 44 writes data from the message buffers 53 to a file system 54 in the cached disk array
5 29. The UxFS layer 44 also reads data from the file system 54 or a file system cache 51
6 and copies the data into the message buffers 53 for transmission to the network clients 23,
7 24, 25.

8 To maintain the file system 54 in a consistent state during concurrent writes to a
9 file, the UxFS layer maintains file system data structures 52 in random access memory of
10 the data mover 26. To enable recovery of the file system 54 to a consistent state after a
11 system crash, the UxFS layer writes file metadata to a log 55 in the cached disk array
12 during the commit of certain write operations to the file system 54.

14 II. Concurrent read and write operations and the allocation of file system blocks.

15 FIG. 3 shows various read and write interfaces in the UxFS layer. These
16 interfaces include a cached read/write interface 61 for accessing the file system cache 51,
17 an uncached multi-threaded write interface 63, and an uncached read interface 64.

18 The cached read/write interface 61 permits reads and writes to the file system
19 cache 51. If data to be accessed does not reside in the cache, it is staged from the file
20 system 54 to the file system cache 51. Data written to the file system cache 51 from the
21 cached read/write interface 61 is written down to the file system cache during a commit
22 operation. The file data is written down first, followed by writing of new file metadata
23 to the log 55 and then writing of the new metadata to the file system 54.

1 The uncached multi-threaded write interface 63 is used for sector-aligned writes
2 to the file system 54. Sectors of data (e.g., 512 byte blocks) are read from the message
3 buffers (53 in FIG. 2) and written directly to the cached disk array 29. For example, each
4 file block is sector aligned and is 8 K bytes in length. When a sector-aligned write
5 occurs, any cache blocks in the file system cache that include the sectors being written to
6 are invalidated. In effect, the uncached multi-threaded write interface 63 commits file
7 data when writing the file data to the file system 54 in storage. The uncached multi-
8 threaded write interface 63 allows multiple concurrent writes to the same file. If a sector-
9 aligned write changes metadata of a file such as file block allocations, then after the data
10 of the file has been written, the new metadata is written to the log 55, and then the new
11 metadata is written to the file system 54. The new metadata includes modifications to the
12 file's inode, any new or modified indirect blocks, and any modified quota reservation.

13 The uncached read interface 64 reads sectors of data directly from the file system
14 54 into the message buffers (53 in FIG. 2). For example, the read request must have a
15 sector aligned offset and specifies a sector count for the amount of data to be read. The
16 data can be read into multiple message buffers in one input/output operation so long as
17 the sectors to be read are in contiguous file system blocks.

18 Typically, the cached read/write interface 61 is used for reading data from read-
19 write files and from any read-only snapshot copies of the read-write files. The uncached
20 write interface 63 is used for sector-aligned writes to read-write files. If the writes are
21 not sector aligned, then the cached read-write interface 61 is used. The uncached read
22 interface 64 is used for sector-aligned reads when there is no advantage to retaining the

1 data in the file system cache 51; for example, when streaming data to a remote copy of a
2 file.

3 FIG. 4 shows various file system data structures 52 associated with a file. A
4 virtual inode (VNODE) 71 represents the file. The virtual inode 71 is linked to an
5 allocation mutex (mutually exclusive lock) 72, a partial block conflict queue 73, a partial
6 write wait queue 74, an input-output (I/O) list 75, a staging queue 76, and preallocation
7 block lists 77. When a file block is preallocated, it is reserved for use in the on-disk file
8 system 54. A preallocated file block can be linked into the in-memory file block
9 structure in the file system cache 51 as maintained by the UxFS layer 44, and later the
10 preallocated file block can become part of the on-disk file system 54 when the
11 preallocated file block is committed to storage. (An example of the file block structure is
12 shown in FIG. 14.) The write threads of the uncached multi-threaded write interface (63
13 in FIG. 3) use the allocation mutex 72 for serializing preallocation of file metadata blocks
14 and commitment of the preallocated metadata blocks. For a Unix-based file, the
15 preallocated metadata blocks include new indirect blocks, which are added to the file
16 when the file is extended. As described below with reference to FIGS. 15 to 16, one or
17 more new indirect blocks may also be added to a read-write file system when processing
18 a client request to write to a data block that is shared between the read-write file system
19 and a read-only snapshot copy of the read-write file system.

20 Preallocation of the file metadata blocks under control of the allocation mutex
21 prevents multiple writers from allocating the same metadata block. The actual data write
22 is done using asynchronous callbacks within the context of the thread, and does not hold

1 any locks. Since writing to the on-disk storage takes the majority of the time, the
2 preallocation method enhances concurrency, while maintaining data integrity.

3 The preallocation method allows concurrent writes to indirect blocks within the
4 same file. Multiple writers can write to the same indirect block tree concurrently without
5 improper replication of the indirect blocks. Two different indirect blocks will not be
6 allocated for replicating the same indirect block. The write threads use the partial block
7 conflict queue 73 and the partial write wait queue 74 to avoid conflict during partial
8 block write operations, as further described below with reference to FIG. 13.

9 The I/O list 75 maps the message buffers (53 in FIG. 2) to data blocks to be
10 written. The write threads use the I/O list 75 to implement byte range locking. The read
11 threads may also use the I/O for byte-range locking. The data blocks, for example, are
12 512 bytes in length providing sector-level granularity for the byte range locking.
13 Alternatively, the data block length is a multiple of the sector size.

14 In order to prevent the log (55 in FIG. 2) from becoming a bottleneck, the
15 preallocated metadata blocks for multiple write threads writing to the file at the same
16 time are committed together under the same logging lock. Committing more than one
17 allocation under one lock increases the throughput. For this purpose, a staging queue 76
18 is allocated and linked to the file virtual inode 71. Preallocation block lists 77 identify
19 the respective preallocated metadata blocks for the write threads writing to the file. The
20 staging queue 76 receives pointers to the preallocation block lists 77 of the write threads
21 waiting for the allocation mutex 72 of the file for commitment of their preallocated
22 metadata blocks. For example, the staging queue 76 is a conventional circular queue, or
23 the preallocation block lists 77 are linked together into a circular list to form the staging

1 queue. There can be multiple files, and each file can have a respective staging queue
2 waiting for commitment of the file's preallocation block lists. A wait list of staging
3 queues 78 identifies the staging queues waiting for service on a first-come, first-served
4 basis.

5 From a client's view, the write operation performed by a write thread in the
6 uncached write interface is a synchronous operation. The write thread does not return an
7 acknowledgement to the client until the write data has been written down to the file
8 system in storage, and the metadata allocation has been committed to storage.

9 FIGS. 5 and 6 show programming in the Common File System (CFS) layer in the
10 network file server for handling a write request from a client. In a first step 81, if the
11 uncached multi-threaded write interface (63 in FIG. 3) is not turned on for the file
12 system, then execution branches to step 82. For example, the uncached interface can be
13 turned on or off per file system as a mount-time option. In step 82, the CFS layer obtains
14 an exclusive lock upon the file, for example by acquiring the allocation mutex (72 in FIG.
15 4) for the file. Then in step 83, the CFS layer writes a specified number of bytes from the
16 source to the file, starting at a specified byte offset, using the cached read/write interface
17 (61 in FIG. 3). The source, for example, is one or more of the message buffers (53).
18 Then in step 84, the CFS layer releases the exclusive lock upon the file, and processing of
19 the write request is finished.

20 In step 81, if the uncached multi-threaded write interface is turned on for the file
21 system, then execution continues to step 85. In step 85, if the write data specified by the
22 write request is not sector aligned (or the data size is not in multiple sectors), then
23 execution branches to step 82. Otherwise, execution continues from step 85 to step 86.

1 In step 86, the CFS layer acquires a shared lock upon the file. The shared lock
2 prevents the CFS layer from obtaining an exclusive lock upon the file for a concurrent
3 write request (e.g., in step 82). However, as described below, the shared lock upon the
4 file does not prohibit write threads in the UxFS layer from acquiring the allocation mutex
5 (72 in FIG. 4) during the preallocation of metadata blocks or during the commitment of
6 the metadata blocks.

7 In step 87, the CFS layer checks the I/O list (75 in FIG. 4) for a conflict. If there
8 is a conflicting data block on the I/O list, then execution waits until the conflicting data
9 block is flushed out of the I/O list. For example, for serializing the writes with prior
10 reads and writes, write access to any blocks being accessed by prior in-progress reads or
11 writes is delayed until these blocks have been accessed by these prior in-progress reads or
12 writes. Moreover, in certain clustered systems in which direct data access to the file in
13 the data storage is shared with other servers or clients, execution may also wait in step 87
14 for range locks to be released by another server or client sharing direct access to the file.
15 After step 87, execution continues to step 88 in FIG. 6.

16 In step 88 of FIG. 6, the CFS layer writes the specified number of bytes from the
17 source to the file, starting at a specified sector offset, using the uncached multi-threaded
18 write interface (63 in FIG. 3). Then in step 89, the CFS layer invalidates any cached
19 entries for the file system blocks that have been written to in the file system cache (51 in
20 FIG. 3). The invalidation occurs after completion of any reads in progress to these file
21 system blocks. In step 90, the CFS layer releases the shared lock upon the file, and
22 processing of the write request is finished.

1 FIG. 7 shows that the parallel read and write architecture can be used to achieve
2 pipelining, since the data write stage does not involve any metadata interactions. The
3 read or write is divided into three steps, namely inode access for reads and writes and
4 preallocation for writes (S1), asynchronous read or write (S2), and inode access for reads
5 and writes and commit for writes (S3). The preallocation in stage S1 is achieved
6 synchronously, and the allocation mutex (72 in FIG. 4) prevents multiple preallocations
7 from occurring simultaneously for the same file.

8 Once the metadata preallocation stage is complete, the asynchronous write (S2) of
9 the data to disk can be handled independently of the metadata preallocation. The
10 asynchronous write (in stage S2) of the data to disk is the longest stage. With pipelining,
11 multiple asynchronous writes can be performed concurrently. This results in an increase
12 in the number of write operations that can be performed in a given time period.

13 The final commit of the allocations (in stage S3) is also achieved synchronously.
14 The allocation mutex (72 in FIG. 4) prevents preallocation for the same file from
15 occurring at the same time as a commit for the same file. However, multiple commits
16 (S3) for the same file may occur simultaneously by gathering the commit requests
17 together and committing them under the same allocation mutex.

18 As shown in FIG. 8, the read and write operations can be pipelined through
19 multiple processors. In this example, a multi-processor board 501 includes two Pentium
20 IV Zeon™ processor chips 502, 503. Each processor chip includes two logical central
21 processing units (CPU) 504, 505 and 506, 507 respectively. Each logical CPU consists
22 of a respective set of on-chip processor registers that share the functional units, input-
23 output ports and cache memory on the chip.

1 The processing of a multi-threaded application by the two logical processors on
2 the same physical processor is called “Hyper-Threading technology.” See, for example,
3 “Building Cutting-Edge Server Applications, Intel® Xeon™ Processor Family Features
4 the Intel NetBurst™ Microarchitecture with Hyper-Threading Technology,” Intel
5 Corporation, 2002, and Chapter 7, “Multiprocessor and Hyper-Threading Technology,”
6 in the Intel® Pentium™ 4 and Intel® Xeon™ Processor Optimization Reference Manual,
7 Order No. 248966-05, Intel Corporation, 2002.

8 For pipelined processing of the multiple concurrent read and write operations as
9 shown in FIG. 7, the metadata management for a file (stages S1 and S3) can be
10 performed by one logical processor, and the asynchronous reads and writes (stage S2) can
11 be performed by another logical processor. Each logical processor executes code threads
12 that are independent of the code threads executed by the other logical processors. For
13 example, as shown in FIG. 8, the first logical CPU 504 of the first processor chip 502
14 performs metadata management for files in a first file system (A:), and the first logical
15 CPU 506 of the second processor chip 503 performs metadata management for files in a
16 second file system (B:). The second logical CPU 504, 507 in each processor chip 502,
17 503 performs asynchronous write operations. For example, the second logical CPU 505
18 of the first processor chip 502 performs asynchronous read and write operations for the
19 first file system (A:), and if the second logical CPU 505 has free processing time, then the
20 second logical CPU 505 of the first processor chip 502 performs asynchronous read and
21 write operations for the second file system (B:). In a similar fashion, the second logical
22 CPU 507 of the second processor chip 502 performs asynchronous read and write
23 operations for the second file system (B:), and if the second logical CPU 507 has free

1 processing time, then the second logical CPU 505 of the first processor chip 502
2 performs asynchronous read and write operations for the first file system (A:).

3 In general, data read requests can be pipelined along with the write requests, and
4 separate processing units can service data read and write requests generated by a primary
5 processor that handles metadata management for each file. The primary processor can
6 hand over an input/output list to a separate secondary processing unit that will then go
7 through the input/output list to perform the transfer of data between the message buffers
8 and cache or disk. For a write operation, the separate secondary processing unit will take
9 the data from the network packets, write it to specified disk locations as requested by the
10 primary, and complete the data write to the disk from the network packets. The pipeline
11 architecture allows a primary processor to do the next write metadata preallocation while
12 other secondary processors are still writing data to disk.

13 When a write I/O request arrives at a primary processor or thread, the request is
14 analyzed and if there are any associated metadata operations (stage S1 in FIG. 7), and the
15 associated metadata operations are executed by the primary processor while the block
16 write I/O is pipelined to another separate secondary processing unit. The secondary
17 processing unit will pipeline multiple block write I/Os (stage S2 in FIG. 7), and will
18 commit the write data to the disk independently of the metadata operation. At the end of
19 the write data commit process, the metadata is committed (stage S3 in FIG. 7) to disk as
20 well. The primary processor is freed to perform additional metadata management
21 operations while the secondary processing unit writes the I/O data to the disk.

22 There could be a pool of secondary logical processing units that execute the write
23 tasks and they can be allocated for additional processing tasks by the primary processor.

1 Some tasks are executed only by the primary processor. For example, a primary
2 processor is allocated to a file system when the data mover is rebooted. The processing
3 of the pipeline is based on the fact that the writes are uncached, and once an
4 asynchronous write is issued to a secondary processor, there is no contingency or locking
5 to the file. If there are any contingencies, then they are solved by the primary processor
6 before the write is issued to the secondary processor.

7 FIG. 9 shows programming in the CFS layer in the network file server for
8 handling a read request from a client concurrent with handling a write request to the same
9 file. For clarity, FIG. 9 omits certain steps for handling read-write interactions. These
10 steps are show in FIG. 10 and further described below.

11 In a first step 91 of FIG. 9, the CFS layer obtains a shared lock upon the file for
12 the read request. In addition, for serializing the read with prior writes, the I/O list for the
13 file (75 in FIG. 4) can be accessed, and read access (in steps 93 or 97) to any blocks
14 being written to by prior in-progress writes is delayed until these blocks have been
15 written to by these prior in-progress writes. In step 92, execution continues to step 93 if
16 the data requested by the read request is found in the file system cache (51 in FIG. 2). In
17 step 93, the data requested by the read request is read from the cache. In step 94, the data
18 read from the cache is added to source message buffers (53 in FIG. 2). In step 95, the
19 shared lock upon the file is released for the write request, and the handling of the read
20 request by the CFS layer is done.

21 In step 92, in the requested data is not found in the cache, then execution branches
22 to step 96. In step 96, the last committed version of the inode for the file is accessed to
23 perform a search for the data block on disk containing the requested data. In step 97, the

1 requested data is read from the data block on disk. In step 98, the data read from the data
2 block on disk is added to the cache. After step 98, execution continues to step 94 to add
3 the data to the source message buffers.

4 FIG. 10 shows the behavior of the server when there are read-write interactions
5 during concurrent access of multiple I/O threads to a single file. In particular, the steps in
6 FIG. 10 occur when a read I/O request accesses blocks to which there is a concurrent
7 ongoing write. Each read and write must obtain a shared lock upon the file. A read
8 acquires the shared lock upon the file in step 91, and a write acquires the shared lock
9 upon the file in step 86.

10 The file system cache (51 in FIG. 2) maintains an index or block map including,
11 for each file system block, an I/O in progress (IOP) flag indicating whether or not a read
12 to the file system block on disk is in progress. During a cached read, if the block map
13 indicates that the block of data to be read is not found in the file system, then execution
14 branches from step 92 to step 510. In step 510, if the IOP flag is set, then execution
15 continues to step 511 to wait for the IOP flag to be cleared. Execution loops back from
16 step 511 to step 92.

17 If in step 510 the IOP flag is not set, then execution branches to step 512 to set the
18 IOP flag and to set a generation count for the block to a value of the present read of the
19 block from disk, and to start the read of the disk. The read of the disk is performed in
20 step 96 to get the committed mapping from the inode, and in step 97 to read data from the
21 block on disk. Once the data from the disk is obtained, the IOP flag and the generation
22 count are checked in step 513. If the IOP flag is set and the generation count is the same

1 as it was in step 512 for the read operation, then in step 98 the data is added to the read
2 cache. After step 98, execution continues to step 94.

3 It is possible that in step 513, the IOP flag will be cleared, or the generation count
4 may be different. The IOP flag can be cleared by a concurrent write operation. For
5 example, after a shared lock upon the file is obtained in step 86 for a concurrent write to
6 the file, data is written in step 515 from message buffers to disk. After step 515, in step
7 516, any cache data for the data block is invalidated, and any IOP flag for the data block
8 is cleared. After step 516, execution continues to step 95 to release the shared lock upon
9 the file. It is possible for a concurrent read to this file block to begin just after the IOP
10 flag for the block is cleared in step 516 but before a previous read has reached step 513.
11 In this case, the previous read will find that the IOP flag is set in step 513 but the
12 generation count will have changed, so that step 98 of adding the data to the cache will be
13 skipped. Step 98 is skipped under these circumstances because the data is current for this
14 previous read operation but stale for subsequent read operations.

15 In short, a write request is serviced by finding partial blocks and creating a partial
16 block list, preallocating metadata blocks for the range of block numbers in the inode that
17 is being written, issuing asynchronous write requests, waiting for completion of the
18 asynchronous write requests, getting a block commit lock, committing the preallocated
19 metadata blocks for the range written to in the inode, releasing the block commit lock,
20 starting asynchronous writes for conflict I/Os, finding the range of blocks in the file
21 system cache to be invalidated, invalidating the file system cache blocks for the block
22 range being committed, if there are active readers, marking the cache range as stale data
23 (for example, via the IOP flag and generation count mechanism shown in FIG. 10).

1 In short, a read request serviced by finding the range of file blocks to be read, and
2 if the blocks are in cache, then reading the data from the cache, and if not, then getting
3 the block commit lock, getting the committed mapping from the inode for the range of
4 blocks to be read, releasing the block commit lock (i.e., the allocation mutex), reading
5 data from disk to the buffer cache and source, and if there are stale blocks in the block
6 range (because a write to the blocks occurred during the read), then invalidating the stale
7 blocks.

8 During a read, IOP flags and generation counts can be used to identify stale
9 blocks. For example, when looking up to see whether data to be read is in cache, missing
10 blocks are marked as IOP (IO in Progress) and the generation count is set to a value
11 associated with this read, and then a read will be started. After completing any reads
12 necessary the blocks that were previously marked as IOP are cleared in one of the
13 following ways: (1) if the slot is cleared, then it's been purged and the just completed
14 read should not be entered and cached; (2) if it's marked as IOP then the generation count
15 is checked: if the generation count is the same as set for this read then data for this read is
16 cached in the slot; otherwise, the data for this read is not cached in this slot but otherwise
17 it can be used to satisfy the read request. During the read process, any concurrent writes
18 are simply allowed to proceed. At the end of the write, the entire range of blocks
19 written are invalidated in the cache. If a cache slot in the range is empty, then it is
20 ignored; otherwise, if the slot had a hint then the slot is cleared, and if the slot was IOP
21 then the IOP flag is cleared and any waiting reads (in step 511 of FIG. 10) are awoken
22 and allowed to proceed.

1 Servicing of the concurrent read and write requests as described above prevents
2 writes to a file from being blocked. There is, however, still an issue of concurrent reads
3 and writes to the same blocks in the same file. For some applications, it is desirable to
4 serialize these reads and writes in so that the data returned by a read operation will
5 indicate that the writes are atomic operations. For example, if each of two write
6 operations write to the same two blocks, a read should not return a first block from the
7 first write operation and a second block from the second write operation. This problem
8 can be solved by looking for conflicting blocks for prior in-progress reads and writes
9 before issuing an asynchronous write operation and by looking for conflicting blocks for
10 prior in-progress writes before issuing an asynchronous read operation, and if a conflict is
11 found, waiting for these prior in-progress conflicting operations to complete before the
12 asynchronous write operation or read operation is issued. This can be done by inspection
13 of the block ranges for prior in-progress writes in the I/O list 75 in FIG. 4. For reads, this
14 would be done in step 91 of FIG. 9. In addition, a read could immediately access non-
15 conflicting blocks in the cache, without waiting for the prior-in-progress writes to
16 complete.

17 Instead of using the I/O list to serialize reads and writes to the same file blocks, a
18 “Write In Progress” (WIP) flag could be added to the file system cache block map. In
19 effect, the WIP flag would be a write lock at the file block level of granularity. Before
20 issuing an asynchronous write operation, during the preallocation stage (S1), the primary
21 processor would set the WIP flags for the file system blocks being written to, unless a
22 WIP flag would already be set, in which case, the write operation would need to wait for
23 completion of the prior conflicting write. The WIP flags would be reset in the

1 asynchronous write stage (S2) after writing to each block. Subsequent writes that
2 encountered a set WIP flag within it's own block range would be required to wait before
3 writing to each block. Likewise attempts to read that encounter a set WIP flag would
4 need to wait until the WIP flag is reset by completion of the conflicting write. If a read
5 operation is accessing blocks being written to by prior in-progress writes, then the read
6 operation should not access these blocks until after they have been written to by the prior
7 in-progress writes. For example, in step 91 of FIG. 9, the I/O list (75 in FIG. 4) or the
8 WIP flags can be accessed to determine the conflicting blocks, before attempting to
9 access these blocks in cache. However, the cache can be accessed immediately for
10 blocks that are not being written to by prior, in-progress writes.

11 FIG. 11 shows a flowchart of a write thread in the UxFS layer (44 in FIG. 2). In a
12 first step 101, the write thread gets the allocation mutex (72 in FIG. 4) for the file. Then
13 in step 102, the write thread preallocates metadata blocks for the block range being
14 written to the file. In step 103, the write thread releases the allocation mutex for the file.

15 In step 104, the write thread issues asynchronous write requests for writing to
16 blocks of the file. For example, a list of callbacks is created. There is one callback for
17 each asynchronous write request consisting of up to 64 K bytes of data from one or more
18 contiguous file system blocks. An I/O list is created for each callback. The
19 asynchronous write requests are issued asynchronously, so multiple asynchronous writes
20 may be in progress concurrently. In step 105, the write thread waits for the asynchronous
21 write requests to complete.

22 In step 106, the write thread gets the allocation mutex for the file. In step 107, the
23 write thread commits the preallocated metadata blocks to the file system in storage. The

1 new metadata for the file including the preallocated metadata blocks is committed by
2 being written to the log (55 in FIG. 3). File system metadata such as the file modification
3 time, however, is not committed in step 107 and is not logged. Instead, file system
4 metadata such as the file modification time is updated at a file system sync time during
5 the flushing of file system inodes. Finally, in step 108, the write thread releases the
6 allocation mutex for the file. This method of preallocating and committing metadata
7 blocks does not need any locking or metadata transactions for re-writing to allocated
8 blocks.

9 FIG. 12 is a more detailed flowchart of steps in the write thread for committing
10 the preallocated metadata. In a first step 111, if there is not a previous commit in
11 progress, then execution continues to step 112. In step 112, the thread gets the allocation
12 mutex for the file. Then in step 113, the thread writes new metadata (identified by the
13 thread's preallocation list) to the log in storage. In step 114, the thread writes the new
14 metadata (identified by the thread's preallocation list) to the file system in storage. In
15 step 115, the thread releases the allocation mutex for the file. Finally, in step 116, the
16 thread returns an acknowledgement of the write operation.

17 In step 111, if there was a previous commit in progress, then the thread inserts a
18 pointer to the threads' preallocation list onto the tail of the staging queue for the file. If
19 the staging queue was empty, then the staging queue is put on the wait list of staging
20 queues (78 in FIG. 4). The thread is suspended, waiting for a callback from servicing of
21 the staging queue. In step 118, the metadata identified by the thread's preallocation list is
22 committed when the staging queue is serviced. The staging queue is serviced by
23 obtaining the allocation mutex for the file, writing the new metadata for all of the

1 preallocation lists on the staging queue to the log in storage, then writing this new
2 metadata to the file system in storage, and then releasing the allocation mutex for the file.
3 Once servicing of the staging queue has committed the new metadata for the thread's
4 preallocation list, execution of the thread is resumed in step 116 to return an
5 acknowledgement of the write operation. After step 116, the thread is finished with the
6 write operation.

7 FIG. 13 is a block diagram of a partial block write during a copy-on-write
8 operation. Such an operation involves copying a portion of the data from an original file
9 system block 121 to a newly allocated file system block 123, and writing a new partial
10 block of data 122 to the newly allocated file system block. The portion of the data from
11 the original file system block becomes merged with the new partial block of data 122. If
12 the new partial block of data is sector aligned, then the partial block write can be
13 performed by the uncached multi-threaded write interface (63 in FIG. 3). Otherwise, if
14 the new partial block of data were not sector aligned, then the partial block write would
15 be performed by the cached read/write interface (61 in FIG. 3).

16 The copy-on-write operation may frequently occur in a file system including one
17 or more read-only file snapshot copies of a read-write file. Such a file system is
18 described in Chutani, Sailesh, et al., "The Episode File System," Carnegie Mellon
19 University IT Center, Pittsburgh, PA, June 1991, incorporated herein by reference. Each
20 read-only snapshot copy is the state of the read-write file at a respective point in time.
21 Read-only snapshot copies can be used for on-line data backup and data mining tasks.

22 In a copy-on-write file versioning method, the read-only snapshot copy initially
23 includes only a copy of the inode of the original file. Therefore the read-only snapshot

1 copy initially shares all of the data blocks as well as any indirect blocks of the original
2 file. When the original file is modified, new blocks are allocated and linked to the
3 original file inode to save the new data, and the original data blocks are retained and
4 linked to the inode of the read-only snapshot copy. The result is that disk space is saved
5 by only saving the difference between two consecutive snapshot copies. This process is
6 shown in FIGS. 13, 14, and 15.

7 FIG. 14 shows a read-write file as maintained by the UxFS layer. The file has a
8 hierarchical organization, depicted as an inverted tree. The file includes a read-write
9 inode 131, a data block 132 and an indirect block 133 linked to the read-write inode, a
10 data block 134 and an indirect block 135 linked to the indirect block 133, and data blocks
11 136 and 137 linked to the indirect block 135.

12 When a read-only snapshot copy of a read-write file is created, a new inode for
13 the read-only snapshot copy is allocated. The read-write file inode and file handle remain
14 the same. After allocation of the new inode, the read-write file is locked and the new
15 inode is populated from the contents of the read-write file inode. Then the read-write file
16 inode itself is modified, the transaction is committed, and the lock on the read-write file
17 is released.

18 The allocation of blocks during the copy-on-write to the read-write file raises the
19 possibility of the supply of free storage being used up after writing to a small fraction of
20 the blocks of the read-write file. To eliminate this possibility, the read-write file can be
21 provided with a “persistent reservation” mechanism so that the creation of a read-only
22 snapshot copy will fail unless there can be reserved a number of free storage blocks equal
23 to the number of blocks that become shared between the read-only snapshot copy and the

1 read-write file. The number of reserved blocks can be maintained as an attribute of the
2 file. The number of reserved blocks for a read-only file can be incremented as blocks
3 become shared with a read-only snapshot copy, and decremented as blocks are allocated
4 during the writes to the read-write file.

5 FIG. 15 shows the read-write file of FIG. 14 after creation of a read-only snapshot
6 copy of the read-write file. The read-only inode 138 is a copy of the read-write inode
7 131. The read-write inode 131 has been modified to indicate that the data block 132 and
8 the indirect block 133 are shared with a read-only snapshot copy. For example, in the
9 read-write inode 131, the most significant bit in each of the pointers to data block 132
10 and the indirect block 133 have been set to indicate that the pointers point to blocks that
11 are shared with the read-write file. (The links represented by such pointers to shared
12 blocks are indicated by dotted lines in FIGS. 15 and 16.) Also, by inheritance, any and
13 all of the descendants of a shared block are also shared blocks. Routines in the UxFS
14 layer that use the pointers to locate the pointed-to file system blocks simply mask out the
15 most significant to determine the block addresses.

16 In general, for the case in which there are multiple versions of a file sharing file
17 blocks, when a file block is shared, it is desirable to designate the oldest snapshot copy
18 sharing the block to be the owner of the block, and any other files to be non-owners of
19 the block. A pointer in a non-shared block pointing to a shared block will have its most
20 significant bit set if the block is not owned by the owner of the non-shared block, and will
21 have its most significant bit clear if the block is owned by the owner of the non-shared
22 block.

1 When writing to a specified sector of a file, a search of the file block hierarchy is
2 done starting with the read-write inode, in order to find the file block containing the
3 specified sector. Upon finding a pointer indicating that the pointed-to block is shared, the
4 pointed-to block and its descendants are noted as “copy on write” blocks. If the specified
5 sector is found in a “copy on write” block, then a new file block is allocated.

6 In practice, multiple write threads are executed concurrently, so that more than
7 one concurrent write thread could determine a need to preallocate the same new file
8 block. The allocation mutex is used to serialize the allocation process so more than one
9 preallocation of a new file block does not occur. For example, once the write thread has
10 obtained the allocation mutex, the write thread then determines whether a new block is
11 needed, and if so, then the write thread preallocates the new block. The write thread may
12 obtain the allocation mutex, allocate multiple new blocks in this fashion, and then release
13 the allocation mutex. For example, to write to a data block of a file, when the write
14 thread finds a shared block on the path in the file hierarchy down to the data block of the
15 file, the write thread obtains the allocation mutex, and then allocates all the shared blocks
16 that it then finds down the path in the file hierarchy down to and including the data block,
17 and then release the allocation mutex.

18 Once a new file block has been allocated, a partial block write to the new file
19 block is performed, unless the write operation writes new data to the entire block. The
20 new file block is the same type (direct or indirect) as the original “copy on write” file
21 block containing the specified sector. If the write operation writes new data to the entire
22 new file block, then no copy need be done and the new data is simply written into the
23 newly allocated block. (A partial write could be performed when the write operation

1 writes new data to the entire block, although this would not provide the best
2 performance.)

3 If the read-write inode or a block owned by the read-write file was a parent of the
4 original "copy on write" block, then the new file block becomes a child of the read-write
5 inode or the block owned by the read-write file. Otherwise, the new file block becomes
6 the child of a newly allocated indirect block. In particular, copies are made of all of the
7 "copy on write" indirect blocks that are descendants of the read-write inode and are also
8 predecessors of the original "copy on write" file block.

9 For example, assume that a write request specifies a sector found to be in the data
10 block 137 of FIG. 15. Upon searching down the hierarchy from the read-write inode 131,
11 it is noted that indirect blocks 133 and 135 and the data block 137 are "copy on write"
12 blocks. As shown in FIG. 16, new indirect blocks 139 and 140 and a new data block 141
13 have been allocated. The new data block 141 is a copy of the original data block 136
14 except that it includes the new data of the write operation. The new indirect block 140 is
15 a copy of the original indirect block 135 except it has a new pointer pointing to the new
16 data block 141 instead of the original data block 137. The new indirect block 139 is a
17 copy of the original indirect block 133 except it has a new pointer pointing to the new
18 indirect block 140 instead of the original indirect block 135. Also, the read-write inode
19 131 has been modified to replace the pointer to the original indirect block 133 with a
20 pointer to the new indirect block 139.

21 In some instances, a write to the read-write file will require the allocation of a
22 new data block without any copying from an original data block. This occurs when there
23 is a full block write, a partial block write to a hole in the file, or a partial block write to an

1 extended portion of a file. When there is a partial block write to a hole in the file or a
2 partial block write to the extended portion of a file, the partial block of new data is
3 written to the newly allocated data block, and the remaining portion of the newly
4 allocated data block is filled in with zero data.

5 It is possible that the UxFS layer will receive multiple concurrent writes that all
6 require new data to be written to the same newly allocated block. These multiple
7 concurrent writes need to be synchronized so that only one new block will be allocated
8 and the later one of the threads will not read old data from the original block and copy the
9 old data onto the new data from an earlier one of the threads. The UxFS layer detects the
10 first such write request and puts a corresponding entry into the partial block conflict
11 queue (73 in FIG. 4). The UxFS layer detects the second such write request, determines
12 that it is conflicting upon inspection of the partial block conflict queue, places an entry to
13 the second such write request in the partial write wait queue (74 in FIG. 4), and suspends
14 the write thread for the second such write request until the conflict is resolved.

15 FIG. 17 is a flowchart of steps in a write thread for performing the partial block
16 write operation of FIG. 13. In a first step 151 of FIG. 17, if the newly allocated file
17 system block (124 in FIG. 13) is not on the partial block conflict queue (73 in FIG. 4),
18 then execution branches to step 152. In step 152, the partial block write thread puts the
19 new block on the partial block conflict queue. In step 153, the partial block write thread
20 copies data that will not be overwritten by the partial block write, the data being copied
21 from the original file system block to the new file system block. In step 154,
22 asynchronous write operations are performed to write the new partial block of data to the
23 new block. In step 155, the partial block write thread gets the allocation mutex for the

1 file, commits the preallocated metadata (or the preallocated metadata is gathered and
2 committed upon servicing of the staging queue if a previous commit is in progress),
3 removes the new block from the partial block conflict queue, issues asynchronous writes
4 for any corresponding blocks on the partial write wait queue, and releases the allocation
5 mutex.

6 In step 151, if the newly allocated file system block was on the partial block
7 conflict queue, then execution continues to step 156. In step 156, the partial block write
8 thread puts a write callback on the partial write wait queue for the file. Then execution is
9 suspended until the callback occurs (from the completion of the asynchronous writes
10 issued in step 155). Upon resuming, in step 157, the partial block write thread gets the
11 allocation mutex for the file, commits the preallocated metadata (or the preallocated
12 metadata is gathered and committed upon servicing of the staging queue if a previous
13 commit is in progress), and releases the allocation mutex.

14 FIG. 18 shows steps in a write thread for allocating file blocks when writing to a
15 file having read-only versions. In a first step 161, if the file block being written to is not
16 shared with a read-only version, then execution branches to step 162 to write directly to
17 the block without any transaction. In other words, there is no need for allocating any
18 additional blocks.

19 In step 161, if the file block being written to is shared with a read-only version,
20 then execution continues to step 163. In step 163, if the file block being written to is an
21 indirect block, then execution branches to step 164. In step 164, a new indirect block is
22 allocated, the original indirect block content is copied to the new indirect block, and the
23 new metadata is written to the new indirect block synchronously. If the block's parent is

1 an indirect block shared with a read-only version, then a new indirect block is allocated
2 for copy-on-write of the new block pointer. Any other valid block pointers in this new
3 indirect block point to shared blocks, and therefore the most significant bit in each of
4 these other valid block pointers should be set (as indicated by the dotted line between the
5 indirect blocks 136 and 140 in FIG. 16). For example, just after the original indirect
6 block content is copied to the new indirect block, the most significant bit is set in all valid
7 block pointers in the new indirect block. As described above with respect to FIG. 16, this
8 copy-on-write may require one or more additional indirect blocks to be allocated (such as
9 indirect block 139 in FIG. 16). For example, the tree of a UxFS file may include up to
10 three levels of indirect blocks. All of the file blocks that need to be allocated can be
11 predetermined so that the allocation mutex for the file can be obtained, all of the new
12 blocks that are needed can be allocated together, and then the allocation mutex for the file
13 can be released.

14 In step 163, if the file block being written to is not an indirect block, then
15 execution continues to step 165. This is the case in which the file block being written to
16 is a data block. In step 165, if the write to the file block is not a partial write, then
17 execution branches to step 166. In step 166, a new data block is allocated and the block
18 of new data is written directly to the new data block. If the original block's parent is an
19 indirect block that is shared with a read-only version, then a new indirect block is
20 allocated for copy-on-write of the new block pointer. As described above with respect to
21 FIG. 16, this copy-on-write may require one or more additional indirect blocks to be
22 allocated.

1 In step 167, for the case of a partial write, execution continues from step 156 to
2 step 167 to use the partial write technique as described above with respect to FIG. 13 and
3 FIG. 17.

4 Various parts of the programming for handling a write thread the UxFS layer have
5 been described above with reference to FIGS. 11 to 18. Following is a listing of the steps
6 in the preferred implementation of this programming.

7 1. The write thread receives a write request specifying the source and
8 destination of the data to be written. The source is specified in terms of message buffers
9 and the message buffer header size. The destination is specified in terms of an offset and
10 number of bytes to be written.

11 2. The write thread calculates the starting and ending logical block number,
12 total block count, and determines whether the starting and ending blocks are partial
13 blocks.

14 3. The write thread gets the allocation mutex for the file.

15 4. The write thread searches the file tree along a path from the file inode to
16 the destination file blocks to determine whether there are any shared blocks along this
17 path. For each such shared block, a new data or indirect block is allocated
18 synchronously, as described above with reference to FIGS. 15, 16, and 18.

19 5. The write thread identifies partial blocks of write data using the starting
20 physical block number and the number of blocks to be written. Only the starting and
21 ending block to be written can be partial. Also, if some other thread got to these blocks
22 first, the block mapping may already exist and the “copy-on-write” will be done by the
23 prior thread. The partial block conflict queue is checked to determine whether such an

1 allocation and “copy-on-write” is being done by a prior thread. If so, the block write of
2 the present thread is added to the partial write wait queue, as described above with
3 reference to FIG. 17.

4 6. The write thread preallocates the metadata blocks.

5 7. The write thread releases the allocation mutex.

6 8. The write threads determine the state of the block write. The block write
7 can be in one of three states, namely:

8 1. Partial, in-progress writes. These are writes to blocks that are on the
9 conflict list. This write is deferred. The information to write out these
10 blocks is added to the partial write wait queue.

11 2. Whole Block Writes.

12 3. Partial, not-in-progress writes. These are partial writes to newly allocated
13 blocks, and are the first write to these blocks.

14 9. The I/O list is split apart if there are any non-contiguous areas to be
15 written.

16 10. Asynchronous write requests are issued for blocks in state 2 (full block
17 writes).

18 11. Synchronous read requests are issued for blocks in state 3 (Partial not-in-
19 progress writes).

20 12. Asynchronous write requests are issued for blocks in state 3.

21 13. The write thread waits for all writes to complete, including the ones in
22 state 1. The write thread waits for all asynchronous write callbacks. The asynchronous
23 writes for blocks in state 1 are actually issued by other threads.

1 14. The write thread gets the allocation mutex.

2 15. The write thread commits the preallocated metadata. The allocation lists
3 being committed are gathered together if a previous commit is in progress, and are
4 written out under the same logging lock as described above with reference to FIG. 12.

5 16. The write thread removes any blocks that the write thread had added to
6 partial block conflict queue, and issues asynchronous writes for corresponding blocks on
7 the partial write wait queue.

8 17. The write thread releases the allocation mutex. The write thread has
9 completed the write operation.

10
11 III. Maintenance of a file version set including read-only and read-write snapshot
12 copies of a production file.

13 As described above with reference to FIGS. 14 to 16 and 18, it is possible to use a
14 copy-on-write technique for creating a read-only snapshot of a Unix-based file. The
15 read-only snapshot can be used for non-disruptive backup by copying the read-only
16 snapshot to a backup media such as magnetic tape or optical disk. In this case the backup
17 is non-disruptive because the backup can be done as a background process while the
18 original read-write file can be accessed on a priority basis. Once a backup copy of the
19 read-only snapshot has been made, then the read-only snapshot can be deleted. For
20 example, the read-only snapshot is deleted by relinquishing the ownership of all of its
21 shared blocks back to the original read-write file, and then de-allocating all of the file
22 system blocks that are exclusively owned by the read-only snapshot.

1 Instead of using a single read-only snapshot for making a backup copy of a file, it
2 is possible to keep a series of read-only snapshots in the network file server. In this case,
3 when a crash occurs and the most recent snapshot is found to be corrupted, then an older
4 snapshot is immediately available for use in restoring the read-write file. Moreover, once
5 an entire copy of an initial snapshot has been migrated to the backup storage, only the
6 changes between the snapshots need be written to the backup storage in order to fully
7 recover all of the snapshots. In this case, there is a savings in backup processing time
8 and in backup storage capacity because more than one backup copy of each file system
9 block will neither be transmitted to the backup storage device nor stored in the backup
10 storage.

11 It is also desirable to provide a non-disruptive and virtually instantaneous
12 mechanism for making a read-write snapshot. For example, during the recovery process,
13 it is often desirable to create a temporary read-write copy of a read-only snapshot prior to
14 restoring the original read-write file after a system crash. Recovery can be attempted
15 upon the temporary read-write file, and then application programs can be tested upon the
16 temporary read-write copy. If a recovery program or an application program should
17 crash when using the temporary read-write copy, then the temporary read-write copy can
18 be deleted, and the recovery process can be restarted using another temporary read-write
19 copy of another read-only snapshot.

20 In order to facilitate the use of multiple read-only and read-write snapshot copies,
21 it is desirable to define a file version set including read-only and read-write snapshot
22 copies produced from an original read-write file. The original read-write file will be
23 referred to as the production file. The read-only snapshot copies will be referred to as

1 read-only versions, or simply versions. The read-write snapshot copies will be referred to
2 as branch files.

3 Shown in FIG. 19 is a preferred logical organization of such a file version set.
4 The file version set includes a production inode 171 for the production file, version
5 inodes 172, 173, 174 for a series of three read-only snapshots of the production file, and
6 two branch inodes 175, 176 for respective read-write copies of the most recent read-only
7 snapshot copy of the production file. The version set also includes a pool 177 of
8 exclusively owned and shared data blocks and indirect file blocks. Each data block or
9 indirect block in the pool 177 is linked to one or more of the inodes 171-176 either
10 directly or indirectly through an indirect block in the pool 177. As will be described
11 below with reference to FIGS. 21 to 23, the inodes 171 to 174 in the version set have a
12 modified format so that the inodes can be linked together via certain inode attributes.

13 Initially, the production file can contain a raw volume of allocated file blocks, or
14 the production file can be a sparse file that has no allocated blocks at creation time. For
15 the case of a sparse file, the initial read-only versions of the production file will be sparse
16 as well. As data is written to a sparse production file, the size of the file can grow up to a
17 pre-specified maximum number of blocks, and the maximum block size can then be
18 extended by moving the end-of-file (eof).

19 As shown in FIG. 20, a new production file is created as either a sparse file or a
20 fully preallocated file. For the case of a sparse file, execution branches from step 331 to
21 step 332 to initially allocate just the inode for the new sparse file. Otherwise, execution
22 continues from step 331 to step 333 to allocate an inode for the new fully preallocated
23 file. Then in step 334, all of the data blocks are allocated for a specified size for the new

1 fully preallocated file. Finally, in step 335, any and all indirect blocks are allocated for
2 the new fully preallocated file as needed to link any of the data blocks of the fully
3 preallocated file to the inode of the fully preallocated file. In other words, a fully
4 preallocated file is created with all of its metadata allocated, including all of its indirect
5 blocks and the data block pointers.

6 By initially allocating all of the metadata for a production file, the overhead
7 associated with the allocations, such as synchronization with concurrent allocations, is
8 eliminated for subsequent writes to the production file. A fully allocated production file
9 provides similar behavior as a storage volume, where all the data blocks are present at the
10 time of creation. A fully allocated production file, for example, is useful as a container
11 for storage objects that are known to be dense, such as video files or copies of raw disk.

12 The initial working file can also be created sparse by writing only to the inode and
13 last block of the file. The sparse file allows the production file to use only those blocks
14 that the client writes data to. This allows less disk blocks to be consumed initially. The
15 sparse file can then be used as the production file for the file version set. Since the new
16 production file after creating a snapshot copy uses new data blocks to write out the data,
17 it results in efficient data block usage, eliminating the need to allocate data blocks that
18 may never be used. The data block allocation scheme can allocate blocks for the new
19 working file in a way that can provide contiguity with the allocated blocks on the
20 previous snapshot copy allowing sequential access to the data blocks for better read
21 performance.

22 For management of the version set of FIG. 19, there is provided a protocol of
23 operations upon the version set. These operations include file creation, file deletion,

1 refresh, and recovery. File creation involves the creation of a read-only snapshot copy
2 from the production file or from a branch file, or the creation of a branch file off a read-
3 only version. File deletion involves the deletion of a read-only snapshot copy or a branch
4 file. Refresh involves discarding the contents of an existing read-only snapshot copy and
5 creating a new snapshot copy using the same name. Restore involves discarding the
6 contents of the production file and creating a new production file using the contents of a
7 specified read-only version.

8 FIG. 21 shows some of the fields of a conventional inode 180. The inode 180
9 includes a mode attribute (MODE) field 181, an access time attribute (ATIME) field 182,
10 an inode change time attribute (CTIME) field 183, one or more data block pointer fields
11 184, and one or more indirect block pointer fields 185.

12 FIG. 22 is a block diagram of an inode 190 in the file version set of FIG. 19. The
13 mode attribute 191 is set with a value IFVERSIONFILE indicating that the inode 190 is
14 for a file version set and the inode has a modified format, as further shown in FIG. 22.
15 The ATIME field 192 in the modified inode 190 stores a version pointer instead of an
16 access time. The CTIME field 193 in the modified inode 190 stores a branch pointer
17 instead of an inode change time. In addition to a data block pointer, the data block
18 pointer field 194 stores a non-owner flag 196 in the most significant bit position. The
19 non-owner flag 196 has a value of zero to indicate that the file is an owner of the data
20 block, and has a value of one to indicate that the file is a non-owner of the data block. In
21 addition to an indirect block pointer, the indirect block pointer field 195 stores a non-
22 owner flag 197 in the most significant bit position. The non-owner flag 197 has a value

1 of zero to indicate that the file is an owner of the indirect block, and has a value of one to
2 indicate that the file is a non-owner of the indirect block.

3 When there is only a production file, with no read-only snapshot copies, the
4 production file owns all of its blocks. When the first read-only snapshot copy file is
5 created, all of the blocks are passed to the new snapshot copy file and it becomes the
6 owner of all of the blocks. The production file still uses the same blocks and the same
7 blocks have identical contents (at least initially); however, it has become a non-owner of
8 those blocks. If any block of the production file is modified, then a new version of that
9 block is allocated and the production file will own that new block. (The new version of
10 the block will be a different block of storage mapped to the same logical address in the
11 file as the original version of the block.) As more snapshot files are created, different
12 snapshot files may own different versions of a block. The owner of any particular block
13 will always be the oldest snapshot copy that uses an identical version of a block, and the
14 oldest snapshot copy will always own all of its blocks. When a sparse file is used, each
15 time a new block is written to it will use the same UxFS allocation mechanism regardless
16 of who owns the data block, the production file or one of the snapshot copies.

17 The concept of a non-owner block is further extended, for indirect blocks, to
18 include the idea of a hierarchy of blocks. For indirect blocks and indirect block trees, if
19 the non-owner flag is set at any level of the tree, then the non-owner state is assumed for
20 all lower-level block pointers. For example, if a pointer to the first level indirect block is
21 marked as non-owner, then all of the data blocks that it points to are assumed to be non-
22 owner, regardless of the state of the non-owner flag in each of the individual block
23 pointer fields.

1 FIG. 23 further shows the use of the version pointers and the branch pointers for
2 linking the inodes 171-176 of the file version set introduced in FIG. 19. FIG. 23 shows
3 that the version pointers are used to form a linked list linking the production file inode
4 171 to all of the version inodes 172, 173, 174. Single links are used in the linked list to
5 conserve space within the conventional inode structure. The versions are linked from
6 most recent to least recent so that a new version inode can be created without modifying
7 other version inodes. The version pointer 201 of the production file 171 includes a most
8 significant bit that is set to indicate that the inode 171 is the inode of the production file.
9 The less significant bits of the version pointer 201 of the production file inode 171
10 contain the inode number of the most recent version if there is any read-only snapshot
11 copy in the version set, and if not, the inode number of the production file inode.

12 For example, in FIG. 23, the version pointer 201 of the production file inode 171
13 includes the inode number 16 of the third read-only snapshot copy inode 174. Each inode
14 172, 173, 174 of a read-only snapshot copy has a version pointer having a most
15 significant bit that is zero and an inode number of the inode of the next most recent read-
16 only version, or in the case of the oldest read-only version, the inode number of the inode
17 171 for the production file. The version pointer 204 of the inode 174 of the third version
18 contains the inode number 15 of the inode 173 of the second version. The version pointer
19 203 of the inode 174 of the second version contains the inode number 13 of the inode 172
20 of the first version. The version pointer 202 of the inode 172 of the first version contains
21 the inode number 10 of the production file inode 171.

22 The branch pointer in each inode has a most significant bit to that is set to indicate
23 the production file inode or a read-only version inode, and that is zero to indicate a

1 branch inode. The less significant bits of the branch pointer contain an inode number.
2 For the production file inode 171 or a read-only version inode 172, 173, 174, if the less
3 significant bits of the branch pointer contain the inode number of the inode, then there are
4 no branch files based on the production file or read-only snapshot copy file, respectively.
5 Otherwise, the less significant bits of the branch pointer in the production inode 171 or
6 version inode 172, 173, 174 include the inode number of the inode of the most recent
7 branch file based on the production file or read-only snapshot copy file, respectively.
8 The less significant bits of the branch pointer in a branch inode contain the inode number
9 of the next most recent branch file based on the same production file or read-only
10 snapshot copy file, or for the oldest branch inode, the inode number of the base
11 production or read-only snapshot copy file. In other words, if there are more than one
12 branch file based on the production file or a read-only version, then the branch pointers
13 are used to form a linked list of branch inodes off the base inode.

14 For example, in FIG. 23, the branch pointer 211 of the production file 171
15 contains the inode number 10 of the production file inode, since there are no branch files
16 based directly on the production file. The branch pointer 212 of the first version inode
17 172 contains the inode number 13 of the first version inode, since there are no branch
18 files based directly on the first read-only version. The branch pointer 213 of the second
19 version inode 173 contains the inode number 15 of the second version inode, since there
20 are no branch files based directly on the second read-only version. The branch pointer
21 214 of the third version inode 174 contains the inode number 18 of the second branch
22 inode 176. The branch pointer 216 of the second branch inode 176 contains the inode

1 number 17 of the first branch inode 175. The branch pointer 215 of the first branch inode
2 175 contains the inode number 16 of the third version inode 174.

3 In practice, it is desirable to prevent a user from creating a branch directly off the
4 production file, since otherwise it would not be possible to recover the branch file after a
5 disruption. The user can always create a read-write copy of the production file by first
6 creating a read-only snapshot copy of the production file and then creating a branch based
7 on the read-only version. If the branch file would be disrupted, then it could be recovered
8 from the read-only version.

9 Because the production file inode serves as an anchor for the snapshot chain, it is
10 desirable to prevent deletion of the production file if there are any snapshot files. The
11 snapshot files should be deleted first.

12 It is also desirable to prevent a read-only snapshot copy from being deleted if
13 there are any branch files based on the read-only version. Typically, any branch files
14 based on the read-only file would be deleted first. Instead of deleting a branch file, it
15 could be converted to a production file and unlinked from the base version, before
16 deletion of the base version. The branch file could be converted to a production file by a
17 background process of copying all blocks that are not owned by the branch file from the
18 base version to newly allocated blocks for the branch file. In the copying process, all of
19 non-owner flags would be cleared.

20 In addition, it is possible to write some changes to a branch file and then create a
21 read-only snapshot copy of the branch file. In this case, the version pointer in the branch
22 inode would contain the inode number of the inode of the read-only version of the branch
23 file. It would also be possible to create branches off this read-only version. In general,

1 the version inodes and the branch inodes could be linked in a hierarchy of version chains
2 and branch chains depending from the production inode 171.

3 Create and delete operations in a version set are synchronized. Further write
4 operations, which may allocate blocks, are synchronized with delete operations. A
5 shared global mutex (a version lock) is used to insure the integrity of the version and
6 branch chains while searching the chains for a file and while modifying the chains. To
7 prevent deadlocks, when concurrent locks are taken on more than one file in a chain, the
8 locking should be done from the head of the chain backwards through the chain. For
9 example, when two successive versions are concurrently locked to delete the earlier
10 version, a lock is first taken on the later version, and then a lock is taken on the earlier
11 version.

12 For each version set, only one create (snap, refresh, restore, etc.) or delete
13 operation may take place at a time. Additional create or delete operations are serialized,
14 because these operations may be changing more than one file in the version set. The
15 create operations are relatively quick and they will hold the global lock for the duration
16 of the operation. Delete operations can take significantly longer. Delete operations are
17 also controlled to prevent multiple delete operations from occurring at the same time.
18 For this purpose, a flag indicating that a delete operation is taking place and a condition
19 variable are maintained in the production file inode.

20 Typically, a Unix-based file system has a file check facility for checking the
21 integrity of the directories and linkages in a file system. This file check facility is
22 extended to recognize that a production file is in a file version set, and once a file version

1 set is found, to check the integrity of the branch and version chains, and to validate the
2 block pointers, the block ownership, and the block counts of the files in the version set.

3 FIG. 24 is a flowchart of a procedure for creating a read-only version of the
4 production file in the file version set of FIG. 19. In a first step 221, a new inode is
5 allocated for the read-only version. Then in step 222, the production file inode is locked.
6 In step 223, the production file inode is copied to the new inode for the version. In step
7 224, the new version inode is updated; for example, the version pointer is updated to link
8 the new version inode into the version chain off the production inode. In step 225, the
9 production file inode is updated; for example, the version pointer is updated to point to
10 the new version inode and the block pointer fields are updated (by setting the most
11 significant bits to set the non-owner flags) to show that the production file is a non-owner
12 of the pointed-to blocks. Then in step 226, the transaction is committed by writing an
13 entry for the new version creation to the log, and writing the production file inode and the
14 new version inode to the file system in storage. Finally, in step 227, the lock on the
15 production file inode is released.

16 FIG. 25 is a flowchart of a procedure for creating a read-write branch off a base
17 version in the file version set of FIG. 19. In a first step 231, a new inode is allocated for
18 the read-write branch. In step 232, the base version inode is locked. Then in step 233,
19 the base version inode is copied to the new inode for the branch. In step 234, the new
20 branch inode is updated; for example, the branch pointer is set to link the new branch
21 inode into the branch chain off the base inode, and the block pointer fields are updated
22 (by setting the non-owner flags in the most significant bits) to indicate that the branch file
23 is a non-owner. In step 235, the base version inode is updated; for example, the branch

1 pointer is set to point to the new branch inode. In step 236, the transaction is committed;
2 for example, by writing an entry into the log indicating the creation of the new read-write
3 branch off the base version, and by writing the new branch inode and the updated base
4 inode to the file system in storage. Finally, in step 237, the lock on the base version
5 inode is released.

6 FIG. 26 shows a procedure for deleting a read-only version in the file version set
7 of FIG. 19, while retaining the next most recent snapshot copy (or the production file,
8 when the snapshot copy being deleted is the most recent read-only version). This
9 involves deleting blocks that are exclusively owned by the snapshot copy being deleted,
10 and retaining blocks that are shared between the snapshot copy being deleted and the next
11 most recent version.

12 In a first step 241 of FIG. 26, a lock is taken on the inode of the read-only
13 snapshot copy and the inode of the next most recent snapshot copy (or the production file
14 if the read-only snapshot copy being delete is the most recent read-only version). The
15 lock prevents the deletion operation from changing the file mapping at the same time that
16 new allocations are being done. If the read-only snapshot copy being deleted is the most
17 recent read-only version, then this lock on the production file is taken in shared mode by
18 writes (and allocations) to prevent blocks owned by the most recent version and not
19 owned by the production file from being passed up to the production file (in step 242) at
20 the same time that new blocks are being allocated.

21 In step 242, there is begun a search for blocks indexed in the inode of the read-
22 only version and corresponding blocks in the inode of the next most recent version (or in
23 the production file if the read-only snapshot copy being deleted is the most recent read-

1 only version). A block in the next most recent version corresponds to a block in the read-
2 only snapshot copy being deleted if the two blocks map to the same range of logical
3 addresses in the two files. The corresponding block may be an identical version of a
4 block (i.e., the same block of storage), in which case the contents will also be the same
5 (because the copy-on-write technique would be used to allocate a new storage block if
6 the contents would change).

7 The search for the corresponding blocks is referred to as a coalescing pass. The
8 objective is to locate blocks that are exclusively owned by the read-only snapshot copy so
9 that these blocks can be freed. Another objective is to locate shared blocks that are
10 owned by the read-only snapshot copy so that ownership of these blocks can be passed to
11 the next read-only snapshot copy (or the production file if the read-only version being
12 deleted is the most recent read-only version). To carry out these objectives, in step 243,
13 the ownership state of each block in the version being deleted is inspected, and a
14 corresponding action is taken depending on the ownership of the block. If the block is
15 not owned by the version being deleted, then an identical version of the block is shared
16 with and owned by an earlier snapshot copy. Also, by inheritance, all of the descendants
17 of the block in the block hierarchy are shared with and owned by an earlier snapshot
18 copy. Therefore, the block (and all of its descendants) can be ignored. The searching
19 process skips over the block and all of its descendants.

20 If the block is owned by the snapshot copy being deleted, then an action is taken
21 depending on the state of the corresponding block in the next most recent version (or the
22 production file if the read-only snapshot copy being deleted is the latest version). If the
23 corresponding block in the next most recent version is not owned, then an identical

1 version of the block is shared between the read-only version being deleted and the next
2 most recent version, and ownership of the block is passed from the read-only version
3 being deleted and the next most recent version. As blocks are passed, the block count is
4 incremented for the next most recent version. If the block being passed is an indirect
5 block, then its descendants become passed by inheritance. However, the indirect block
6 (and any indirect block descendants) should be traversed to count the number of
7 descendants in order to increment the block count for the next most recent snapshot copy
8 by the number of descendants. A function is provided to do the counting for one indirect
9 block, and this function may be called recursively for second and third level indirect
10 trees.

11 If the block is owned by the read-only file version being deleted and the
12 corresponding block in the next most recent file snapshot is owned by the next most
13 recent version, then the block was modified between the read-only snapshot being deleted
14 and the next most recent version. In this case, the read-only snapshot copy being deleted
15 has exclusive ownership of its version of the block, and its version of the block can be
16 freed. If the block is not found in the next most recent version (for example because the
17 extent of the file had been reduced), then the read-only snapshot copy being deleted has
18 exclusive ownership of the block, and the block can be freed.

19 Finally, in step 244, when the search for blocks has been completed, the inode of
20 the read-only snapshot copy being deleted is deallocated, and the lock is released.

21 The deletion of blocks from the read-only snapshot copy being deleted can be
22 done in such a way that truncation occurs from the end of the file backwards. In this
23 case, the file size can be used as a processing indicator, and the deletion process can be

1 halted and restarted. The coalescing and cleanup of the file can be done asynchronously,
2 although only one file deletion from the version set will be performed at any given time.

3 In a preferred implementation, the coalescing and cleanup of a file is done by a
4 program loop that executes a series of transactions. Each pass through the program loop
5 executes one transaction. Each transaction is logged, so the coalescing and cleanup can
6 be resumed if interrupted by a system crash. During each transaction, an exclusive lock
7 is held on the next most recent version (or the production file, if the snapshot copy being
8 deleted is the most recent read-only version). This prevents any attempt to allocate
9 blocks in the locked file. The exclusive lock is released at the end of processing for each
10 transaction, in order for any conflicting processes to make forward progress.

11 The process of deleting versions can be simplified when all of the files in the
12 version set are deleted. In this case, all of the blocks in the version set are deallocated. In
13 addition, the deletion of multiple successive versions can be optimized. Only a single
14 coalescing pass is needed to pass blocks that are owned by the successive versions being
15 deleted but shared with the next most recent version being retained. There is no need to
16 pass blocks between two successive versions that will both be deleted.

17 There is no need for passing blocks when a branch file is deleted. Any blocks that
18 are owned by the branch are deallocated, and any non-owner blocks are ignored.

19 FIGS. 27-28 show details of the persistent reservation mechanism ensuring that
20 the sharing of file blocks among the files in the version set of FIG. 19 is not likely to
21 result in a shortage of file blocks when writing to the production file or a branch file. In a
22 first step 251, a number of free blocks are reserved for each read-write file. The number
23 is maintained as a "block reservation" attribute for the read-write file. In step 252, the

1 number of blocks in each file is maintained as a “block count” attribute for the read-write
2 file. In step 253, when a new block is allocated to the read-write file from the block
3 reservation for the file, the block reservation is decremented, and the block count is
4 incremented. Additional free blocks are reserved to prevent the block reservation from
5 becoming negative, or else the allocation fails. In step 254, when a block is removed
6 from the read-write file, then an additional block can be reserved for the file. In this case,
7 the block reservation for the file is incremented, and the block count for the file is
8 decremented.

9 Continuing in step 255 of FIG. 28, the creation of a read-only snapshot copy of a
10 read-write file will fail unless there can be reserved a number of free blocks equal to the
11 block count of the read-write file. For example, more free blocks are reserved as the
12 block reservation count of the read-write file is incremented by the number of blocks that
13 become shared with the new read-only file.

14 In step 256, the creation of a read-write branch of a read-only base snapshot copy
15 will fail unless there can be reserved a number of free blocks equal to the block count of
16 the read-only base version. For example, more free blocks are reserved as the block
17 reservation of the branch file is incremented by the number of blocks that become shared
18 with the new read-write branch file.

19 In step 267, a restore of the production file with a read-only snapshot copy will
20 fail if the block count of the read-only snapshot copy exceeds the block count of the
21 production file unless there can be reserved a number of free blocks equal to the
22 difference between the block count of the read-only snapshot copy and the block count of
23 the production file. For example, the block reservation of the production file is

1 incremented by the original block count of the production file, decremented by the block
2 count of the read-only version, and any deficiency is made up by incrementing the block
3 reservation as additional free blocks are reserved for the production file.

4 FIG. 29 is a state diagram for the process of restoring a production file with a
5 read-only version. The state diagram has an initial state 261 of the original production
6 file, an intermediate state 262 in which the version set has been prepared for a restore
7 operation, and a final state 263 in which the production file has been restored. The
8 process of restoring the production file is provided with a distinct intermediate state
9 because it is possible that the restore operation may fail or it may be desirable to provide
10 the user with an option to abort the restoration process, for example, because sufficient
11 free file system blocks are not available to satisfy the persistent reservation requirement.
12 Therefore, the restoration process has been configured for a two-phase commit process,
13 in which the first phase is to prepare for a restore operation, and the second phase is to
14 either abort the restore operation or commit the restore operation.

15 Once a process capable of failure has been configured for such a two-phase
16 commit process, then it can be used in the well-known two-phase distributed commitment
17 protocol. In the two-phase distributed commitment protocol, the preparation and
18 commitment can be done at multiple distributed sites under management of a single
19 controller. In the first phase, the preparation at all of the sites is performed at the request
20 of the controller, and the results are reported back to the controller. If all sites report
21 back that the preparation has been successful, then the controller may request all of the
22 sites to commit to completing the process. In this case, it is highly probable that the
23 process will be completed everywhere. However, if any one of the sites reports back that

1 its preparation has been unsuccessful, then the controller may request all of the sites to
2 abort their preparation.

3 For restoring files, the two-phase distributed commitment protocol could be
4 useful for preparing to restore multiple files in a distributed data storage system. The
5 files could be distributed across a network and stored in different network file servers. If
6 the preparation for restoration of all of the files would be successful, then the restoration
7 of all of the files would be committed. If the preparation for restoration of any of the
8 files would be unsuccessful, then the restoration of all of the files would be aborted. The
9 preparation for the restoration process could ensure, to a high probability, that all of the
10 files in the file system could be restored together, or else none of them would be restored.

11 FIG. 30 shows a procedure for preparing for the restoration of the production file.
12 In a first step 271, a branch file copy is created from a specified base version. The base
13 version is the read-only snapshot copy to be used for restoring the production file. Also
14 an attempt is made to reserve the difference between the block count of the specified base
15 version and the block count of the production file. In step 272, if the restoration has been
16 prepared, then execution returns reporting success. Otherwise, execution returns
17 reporting failure. For example, execution could return with a fatal error if the specified
18 base version has been corrupted so that no branch file copy could be created. Execution
19 could also return with an indication that creation of the branch file copy was successful
20 but there were insufficient resources for persistent reservation.

21 FIG. 31 shows a procedure for aborting the restoration of the production file. In
22 step 281, the new branch file (created during preparation for the restore) is discarded.
23 Read-write access may continue with the original production file.

1 FIG. 32 shows a procedure for committing the restoration of the production file.
2 In step 291, the new branch file (created during the preparation for the restore) assumes
3 the identity of the production file. This is done by unlinking the branch file inode from
4 the branch chain off the base version inode, linking the branch file inode into the version
5 chain in lieu of the production file inode, and changing the pointer in the parent directory
6 of the production file to point to the branch file inode in lieu of the production file inode.
7 Then the old production file inode and the blocks owned by the old production file are
8 deallocated. Unless a nondestructive restore option is selected, any read-only versions
9 more recent than the base version are also deleted by deallocating all of their owned
10 blocks and then deallocating their inodes.

11 A refresh of a read-only snapshot copy takes an existing version file, discards it
12 contents, and creates a new version for the snapshot file. The new version is a snapshot
13 copy of the present state of the production file. FIG. 33 shows a preferred procedure. In
14 step 301, a new inode is created, and the contents of the original version inode are copied
15 into the new inode. In step 302, the new inode is linked into the version chain in lieu of
16 the original version inode. In step 303, the original version inode is used to create a new
17 snapshot of the production file. In other words, the production file inode is copied to the
18 original version inode, the original version inode is linked into the version chain as the
19 most recent version, and the non-owner flags are set in the production file inode. Then in
20 step 304, the old read-only snapshot copy of the new inode is scheduled for asynchronous
21 deletion. In this fashion, the refreshed snapshot can become available for user access
22 before the old snapshot copy is deleted.

1 It is desirable to provide users with a convenient method of referencing the
2 various files in a version set. A preferred method is to provide a hierarchical naming
3 convention similar to a hierarchical path name common for Unix-based file; for example,
4 a path name for a Unix-based file is typically in the form of:

5 DirectoryName\SubDirectoryName\...\FileName.

6 For referencing files in a version set, a suitable hierarchical naming convention is in the
7 form of:

8 ProductionFileName [\$VersionName][\$\$BranchName]....

9 In other words, a single occurrence of the "\$" symbol is used as a delimiter to indicate a
10 following version name, and a double occurrence of the "\$\$" symbol is used as a
11 delimiter to indicate a following branch name. In accordance with this convention, the
12 six files in the version set of FIG. 19 could have the following file names:

13
14 Production Inode 171: ProductionFileName
15 Version 1 Inode 172: ProductionFileName\$1
16 Version 2 Inode 173: ProductionFileName\$2
17 Version 1 Inode 174: ProductionFileName\$3
18 Branch 1 Inode 175: ProductionFileName\$3\$\$1
19 Branch 2 Inode 176: ProductionFileName\$3\$\$2

20

21 This naming convention would have the advantage that all of the files in the
22 version set could share the same NFS file handle or CIFS file id. In addition, the naming
23 convention would have the advantage that a file name matching the pattern could trigger

1 the creation of a new snapshot copy or branch file. For example, if a request to create a
2 new version specified an existing production file name followed by the delimiter "\$"
3 followed by a version name that did not exist, then a new snapshot of the production file
4 would be created having the specified version name. The file handle returned would be
5 that of the production file.

6 FIGS. 34 and 35 show a procedure for parsing a file name in accordance with the
7 above convention. In a first step 311, version chain scanning is set to begin at the
8 production inode. Then in step 312, the production file name is parsed from the name of
9 the file in the version set. In step 313, if an end of input is reached in the parsing of name
10 of the file in the version set, then execution returns indicating that the production file is to
11 be accessed. Otherwise, execution continues to step 314 to get the next character from
12 the file name input. In step 315, if this next character is not the "\$" character, then
13 execution returns reporting a format error. Otherwise, execution continues to step 316; to
14 parse a version name X and scan the version chain until the inode is found for the version
15 named X. In step 317, if an end of input is reached in the parsing of the version number,
16 then execution returns indicating that the read-only version X of the production file
17 system is to be accessed. Otherwise, execution continues to step 318 in FIG. 35.

18 In step 318 of FIG. 35, the next two characters are obtained from the input of the
19 name of the file in the version set. In step 319, the next two characters are not "\$\$", then
20 execution returns reporting a format error. Otherwise, execution continues to step 320 to
21 parse a branch name Y and scan the branch chain off the version named X until the
22 branch named Y is found. In step 321, if an end of input of the name of the file in the
23 version set has been reached, then execution returns indicating that the branch Y off the

1 read-only snapshot copy X is to be accessed. Otherwise, execution continues to step 322.
2 In step 322, the next character is obtained from the input of the name of the file in the
3 version set. In step 323, if the next character is not "\$", then execution returns reporting
4 a format error. Otherwise, execution continues to step 324 to set the version chain
5 scanning to begin at the inode of branch Y of version X. After step 324, execution loops
6 back to step 316 of FIG. 34.

7 An alternative naming convention could use a directory for the version set. The
8 directory could have an entry for each file in the version set, and an arbitrary name could
9 be assigned to each file in the version set. The directory for the version set could provide
10 a means for locating a branch file that would become unlinked from its base version
11 when its base version is deleted, or locating versions that might be retained after deletion
12 of the production file. This alternative, however, involves additional processing overhead
13 for maintaining the directory entries and keeping track of the directory itself.

14 Another alternative is to use pseudo directories. Each pseudo directory could
15 have a version date or user supplied label associated with it. Also, it could have a
16 specific file system version level number. Reading the pseudo directory could return a
17 list the files that had a version number less than or equal to the version number of the
18 pseudo directory. This has an advantage in that it is somewhat easier to manage older file
19 versions, since they are collected together in the pseudo directories. This alternative
20 would require the production file to exist as a name anchor and would also involve
21 additional processing time for maintaining the pseudo directories.

22 In view of the above, there has been described a way of creating read-only and
23 read-write snapshot copies of a production file in a Unix-based file system. The

1 production file and the snapshot copies are organized as a version set of file inodes and
2 file blocks including blocks that are shared among the snapshots and the production file.
3 The inodes in the version set are linked together by version pointers and branch pointers.
4 The user is able to choose only those production files and snapshots that are considered
5 important enough to be copied and saved. This has the advantage of improving both
6 performance and storage efficiency. A protocol is provided for creating read-only and
7 read-write versions, deleting read-only and read-write versions, restoring the production
8 version with a specified version, refreshing a specified version, and naming the files in
9 the version set. The production file can be created as a fully pre-allocated file by pre-
10 allocation of all of its blocks in the file system at creation time, or as a sparse file whose
11 inode is allocated at creation time and whose other blocks are allocated as needed when
12 its data blocks are written to.

13 When writing to a file block that is shared between the production file and a read-
14 only version, a new block is allocated to the production file. The contents of the shared
15 block are written to the new block if there is a partial write to the new block. This copy-
16 on-write technique is complicated by the presence of indirect blocks, which may also
17 need to be copied. To solve this problem, block pointers are marked with a flag
18 indicating whether or not the pointed-to block is owned by the parent inode. A non-
19 owner marking is inherited by all of the descendants of a block. The block ownership
20 controls the copying of indirect blocks when writing to the production file, and also
21 controls deallocation and passing of blocks when deleting a specified read-only version.

22